We claim:

1. A modular computing system comprising :

a set of functionally independent processing nodes having one or more local processors and local memory, wherein each processing node includes a common communication interface for communicating with other nodes within the system via a messages conforming to a packetized network protocol; and

one or more routing modules communicatively coupling the processing nodes via their respective common communication interface.

2. The system of claim 1, wherein the nodes operate in a global shared-memory address space.

3. The system of claim 1, wherein the common communication interface within each node connects the local memory to the local processors, provides at least one port for interfacing with the routing modules and at least one port for communicating with an input/output (I/O) subsystem.

4. The system of claim 1, wherein the common communication interface of each node may be directly coupled together, thereby eliminating the routing module.

5. The system of claim 1, wherein the computing system includes a system control hierarchy comprising a level one controller within each node, a level two controller providing rack-wide control and a level three controller providing system-wide control.

6. The system of claim 5, wherein the level one controller within each node controls direct low-level communications within the node.

7. The system of claim 5, wherein each routing module comprises a level two controller

Cont A1

8. The system of claim 5, wherein the level three controller is a standalone workstation.

9. The system of claim 1, wherein the nodes are communicatively coupled to the routing modules via a high-speed Universal Serial Bus.

10. The system of claim 1, wherein each common communication interface includes:

a processor interface for interfacing to one or more processing nodes.

a memory interface for interfacing to local memory as a portion of the global memory and for maintaining cache coherency across the computing system.

an I/O interface for communicating with an I/O subsystem.

11. The system of claim 1, wherein the common communication interface includes a plurality of interface control units.

12. The system of claim 11, wherein the common communication interface includes a central crossbar communicatively coupling each interface control unit for the exchange of data between the external interfaces at high data rates.

13. The system of claim 12, wherein each interface control unit within common communication interface communicates by sending messages through the crossbar.

14. The system of claim 1, wherein the message protocol is a synchronous message protocol comprising requests and replies.

15. The system of claim 12, wherein the crossbar coverts the messages to an internal message format.

cont A1

16. The system of claim 12, wherein crossbar communicates the messages across two internal two virtual channels by multiplexing the messages across physical channels connecting each unit.

17. A processing node for a modular computing system comprising :

one or more local processors;

local memory;

a common communication interface coupled to the local processors and the local memory, wherein the common communication interface includes:

a processor interface for communicating with one or more external processing nodes;

a memory interface by which the local processors and the external processor nodes communicating with the local memory;

a routing interface for communicating with an external routing module; and

an I/O interface for communicating with an I/O external subsystem.

18. The processing node of claim 17, wherein the nodes operate in a global shared-memory address space.

19. The processing node of claim 17, wherein the common communication interface of the node may be directly coupled to a common communication interface of another such node via the I/O interface.

20. The processing node of claim 5 and further including a system controller for directing low-level communications within the node.

21. The processing node of claim 17, wherein the routing interface includes a high-speed Universal Serial Bus.

Cont A1

22. The processing node of claim 17, wherein the common communication interface includes a plurality of interface control units.

23. The processing node of claim 22, wherein the common communication interface includes a central crossbar communicatively coupling each interface control unit for the exchange of data between the external interfaces at high data rates.

24. The processing node of claim 22, wherein each interface control unit within common communication interface communicates by sending messages through the crossbar.

25. The processing node of claim 17, wherein the message protocol is a synchronous message protocol comprising requests and replies.

26. The processing node of claim 22, wherein the crossbar coverts the messages to an internal message format.

27. The system of claim 22, wherein crossbar communicates the messages across two internal two virtual channels by multiplexing the messages across physical channels connecting each unit.

28. A modular computing system comprising :

a set of functionally independent processing nodes operating in a global, shared address space, wherein each node has one or more local processors and local memory, wherein each processing node includes a common communication interface for communicating with other modules within the system via a message protocol, and further wherein the common communication interface provides a single high-speed communications center within each node to operatively couple the node to one or more external processing nodes, an external routing module, or an input/output (I/O) module.

Cont A1

29. The system of claim 28, wherein the computing system may include an arbitrary combination of processing nodes and other modules such that there need not be a fixed relation between the number of processing nodes and the other modules.

30. A modular computing system comprising :

a set of functionally independent processing nodes that can be operatively coupled to form one of a plurality of computing topologies, wherein each computing topology supports a number of the processing nodes, and further wherein each computing topology is a superset of the computing topologies that support fewer processing nodes.

31. The computing system of claim 28, wherein each processing node includes a common communication interface for communicating with other processing nodes within the system